

PTO/PCT Rec'd 04 MAY 2000

- 1 -

Improvements in, or relating to, ATM Telecommunication Systems

5 The present invention relates to an ATM telecommunication system adapted for the transmission of IP data over ATM, and to a simple IP/ATM method (SIAM) for the transportation of IP data over an ATM transmission system.

IP over ATM

10 The Internet Engineering Task Force (IETF) has proposed different methods for the transportation of IP (Internet Protocol) over ATM (Asynchronous Transfer Mode). For example, Classical IP over ATM is the name of a protocol supporting automatic address resolution of IP addresses (see M. Laubach, Classical IP and ARP over ATM, RFC 1577, January 1994). This protocol introduces the notion of LIS (Logical IP Subnet) consisting of a group of IP nodes (IP hosts, or routers), members of the same IP subnet and connected to a single ATM network. Each LIS supports a single ATM ARP server (Address Resolution Server) that resolves the IP addresses of the LIS. Any packet for a destination outside the LIS is sent to a default router. This approach is very simple but has a significant drawback because it needs a conventional router to interconnect different LISs although the source and destination attach to the same ATM network. This hop-by-hop routing approach implies throughput and latency problems.

20 At the present time, IETF is working on a new protocol called NHRP (Next Hop Resolution Protocol) that solves the hop-by-hop problem by allowing short-cut routing. In place of ARP servers, NHRP servers (NHSs) are used where each maintains the IP to ATM address mappings of all nodes associated with that NHS as well as the learned mappings. When a host needs to transmit a packet across the ATM network, it resolves the IP destination address by asking its NHS which resolves the address, or forwards the resolution request to the next NHS on the path to the destination. When the destination address is resolved, the sender sets up a direct SVC (Switched Virtual Connection) to that destination so that short-cut routing is achieved. One significant drawback of the NHRP approach is that looping may occur in some cases. Furthermore, it scales badly because no aggregation of user data flows is supported, leading to the consumption of a large number of SVCs.

30 At the present time, the ATM Forum is working on MPOA (Multi-Protocol Over

- 2 -

ATM) for the transport of IP and other Internet-working protocols over ATM (see ATM Forum, Multi-Protocol over ATM, straw ballot, str-mpoa-mpoa-01.00, February 1997). MPOA supports virtual private networking regardless of physical location. In addition, MPOA:

5

- defines a high performance and low-latency method to support IP across an ATM network based on short-cut routing; and
- makes use of LAN emulation for intra-subnet communication and NHRP for intersubnet communication; but
- suffers from the same drawbacks as NHRP.

MPOA is also very complex and requires the use of ATM Forum's LAN emulation at the MPOA clients.

A number of industrial IP/ATM approaches have recently been defined. Some examples of these approaches are:

20

- IP switching from Ipsilon (see P Newman et al, Transmission of Flow Labelled Ipv4 on ATM Data Links, IETF RFC 1954, May 1996);
- Tag Switching from Cisco (see Y Rekhter et al, tag Switching Architecture Extensions for ATM: Overview, <draft-rekhter-tagswitch-arch-00.txt>, January 1997); and
- Cell Switch Router from Toshiba (see Y Katsube et al, Toshiba's Router Architecture Extensions for ATM: Overview, IETF RFC 2098, February 1997).

25

30

These approaches combine the speed and simplicity of ATM with the scalability and control of the IP layer. The major difference between these approaches, compared to existing ones, lies in the degree to which IP is integrated with ATM.

- 3 -

Two major changes in the Internet have also contributed to the birth of new paradigms. The first one is the knowledge of the nature of the applications at the IP level, and the second one is the relaxation of the connectionless nature of IP.

5 A common feature of all the aforementioned approaches is the use of IP flows which are classified and mapped into ATM virtual channels (labelling) achieving high performance and low latency through bypassing the hop-by-hop IP processing. Also, each of the aforementioned approaches makes use of a control protocol that performs labelling and announces the result of that to its neighbouring nodes (label distribution). However, SIAM has a simpler architecture because it does not require the use of the label distribution protocols which characterise the existing methods, for example, IP Switching (Ipsilon Flow Management Protocol), Cell Switch Router (Flow Attribute Notification Protocol) and Tag Switching (Tag Distribution Protocol (Cisco)).

10 It is an object of the present invention to provide a simple IP/ATM method (SIAM) that bypasses hop-by-hop IP processing by mapping IP flows into ATM VCs (Virtual Circuits).

15
20 According to one aspect of the present invention, there is provided, an ATM transmission system, adapted for the transmission of IP data and including a core network of ATM switches, said ATM transmission system being adapted to handle inter-subnet communications, characterised in that said core network includes access IP/ATM nodes (AINs) and core IP/ATM nodes (CINs) for handling said inter-subnet communications, in that each AIN is attached to an ATM switch of the core network through an ATM User-Network Interface (UNI) and is adapted to perform IP data flow classification and labelling for facilitating mapping of IP data flows to ATM VCs, and to communicate with IP/ATM hosts and routers, in that each CIN is attached to an ATM switch of the core network through an ATM UNI and is adapted to perform routing and labelling, in that said AINs and CINs are interconnected through Virtual Path Connections (VCPs), and in that permanent VCPs are set-up between adjacent CINs. The AINs may be adapted to communicate with non-ATM hosts using, for example, Ethernet. The inter-subnet communications may be effected on a hop-by-hop basis, in which case, the ATM transmission system may be adapted to map each IP data flow into

25
30

- 4 -

an ATM Virtual Circuit (VC) for each hop, between nodes, on the communication path towards a destination subnet. In which case, the ATM VC for each hop between nodes will be a VCI selected from a VPC connecting the two nodes.

5 Each AIN may be adapted, on receipt of an IP data packet from an IP/ATM host for transmission to a destination subnet, to reassemble the IP data packet, decrement a TTL (Time To Live) timer, discarding the IP data packet if the timer reaches zero, compute the IP data packet header checksum, and check whether the destination subnet is supported by the AIN and, if so, send the IP data packet on that subnet. Each AIN may include, and be adapted to maintain, a labelling table and a standard IP forwarding table, the entries of which are completed by an IP routing protocol. With this arrangement, each entry in the labelling table includes a subnet address, the outgoing VC (VPI/VCI) to reach the subnet, and a timer for the VC, and the timer is adapted to be started whenever IP packets are sent on the VC. In the event that the destination subnet is not supported by the AIN, it may be adapted to identify an existing ATM VC towards said destination subset and to send the IP data packet on an identified existing ATM VC. In the event that an existing ATM VC is not identified, the AIN may be adapted to identify the next hop on the path to the destination subnet, by consulting its forwarding table, and to chose, and send the IP packet on, a free VCI from the VPC to the next hop.

20 In which case, the AIN is adapted to update its labelling table by making an entry containing the destination subnet address and the chosen VC, and to restart the entry's timer, said timer having a lifetime value for the VC. The next hop on the path to the destination subnet may be a CIN.

25 Each AIN may be adapted to listen to the VPCs to which it is set up, to reassemble the cells coming from the core network into IP packets, to perform computations on TTL and header checksum, and to route valid packets towards ATM hosts/routers attached to an access network.

30 Each CIN may be adapted to listen to the VPCs to which it is set up, retrieve an IP packet from a cell received on a VCI of one of said VPCs, decrement a TTL and compute the IP packet header checksum, consult its labelling table to determine if there is an outgoing VC already in existence to a destination subnet to which the IP packet is

- 5 -

5 to be sent and, if an existing outgoing VC is identified, send the IP packet on that VC and restart the VC's timer. In the event that an existing outgoing VC is not identified, said CIN may be adapted to identify the next hop on the path towards the destination subnet by consulting its IP forwarding table, and to chose a free VCI from the VPC to the next hop. In which case, the CIN is adapted to update its labelling table by making an entry containing the destination subnet address, the incoming VC, the outgoing VC and the free VCI chosen from the VPC to the next hop, and to set the VC timer for the newly created entry in said labelling table, send the IP packet on the outgoing VC and cross-connect the VC incoming to, and the VC outgoing from, its ATM switch. Thereafter, all cells received on said incoming VC are forwarded to the outgoing VC within said ATM switch. The CIN may be adapted to control said cross-connection using Switchlets, an Ariel interface being provided between said ATM switch and said CIN for cross-connecting said incoming and outgoing VCs of said ATM switch. The cross-connected VCs may be the VCIs of the incoming and outgoing Virtual Path (VP) links between said CIN's ATM switch and its two neighbouring ATM switches, one on an input side of, and the other on an output side of, said CIN's ATM switch.

20 The CIN may be adapted to periodically monitor the cross-connected VCs in order to disconnect any VC found to be idle. An Ariel interface may be provided between said CIN and its ATM switch, in which case, the CIN would be adapted to use the Ariel interface to periodically monitor the cross connected VCs.

25 In the event that an incoming VC is found to be inactive, said CIN may be adapted to purge the corresponding outgoing VC from its labelling table, the corresponding incoming VC being cross-connected back to the CIN via a VP link on the CIN/ATM switch interface, and to re-enter the purged outgoing VC into its labelling table as a free VC.

30 Each entry in a CINs labelling table may contains a list of all incoming, merged, Vcs, the outgoing VC towards a destination subnet, said outgoing VC being used by the CIN during the merging of new incoming VCs and is established when a new entry is created in the labelling table, the destination subnet address, and VC timers, one for each incoming VC.

- 6 -

5 Each CIN may be adapted to merge a new incoming VC with an outgoing VC, in which case, each CIN may be adapted to listen to the VPCs to which it is set up and to retrieve an IP packet received on a VCI of one of said VPCs, decrement the TTL and compute the IP packet header checksum, consult its labelling table to determine if there is an outgoing VC already in existence to the destination subnet to which the retrieved IP packet is to be sent, if an existing outgoing VC is identified, send the IP packet on that VC to the destination subnet, the CIN merging the incoming VC to the outgoing VC to the subnet and, in the event that the incoming VC is not included in the list of incoming VCs in its labelling table, add it to said labelling table. In the event that an existing outgoing VC is not identified, each CIN may be adapted to find the next hop on the path towards the destination subnet by consulting its IP forwarding table and to chose a free VCI from the VPC to the next hop. In which case, said CIN may be adapted to update its labelling table by making an entry containing the destination subnet address, the incoming VC (first in the VC list), the outgoing VC and the free VCI chosen from the VPC to the next hop, and to send the IP packet on the outgoing VC, to merge the incoming VC to the outgoing VCs using a CIN/ATM switch interface and to start a timer for the incoming VC. Thereafter, all cells received on said incoming VC are forwarded to the outgoing VC within said ATM switch.

20 Each CIN may be adapted to periodically monitor the cross-connected/merged VCs in order to disconnect any VC found to be idle. In the event that an incoming VC is found to be inactive, said CIN may be adapted to purge the corresponding outgoing VC from its labelling table, the incoming VC being cross-connected back to the CIN via a VP link on the CIN/ATM switch interface. In the event that the VC list becomes empty, said CIN may be adapted to return the purged outgoing VC to its labelling table, and to delete the entry for the subnet from the labelling table.

25 In the event of rerouting of an IP packet, the CIN may be adapted to assign new VCs to the affected subnet addresses in its labelling table, based on new routes for the IP packet, the old VCs, after being timed out, being re-entered in the labelling table as free VCs.

30

- 7 -

The ATM transmission system of the present invention may be adapted to provide Quality of Service (QoS) support by setting up CBR VPCs in parallel to UBR VPCs. With this arrangement, each AIN may be adapted to assign service classes to data flows and give an indication of the data rate assigned to each flow using the type of service field in the IP packet header, and each CIN may be adapted, on receipt of an IP packet on a CBR VC, to first find a free outgoing CBR VC towards the destination subnet, allocate a required data rate based on the type of service value, send the IP packet on the free outgoing CBR VC, and finally, cross-connect the incoming and outgoing VCs at its ATM switch.

According to another aspect of the present invention, there is provided, a method of transmitting IP data over an ATM transmission system having a core network including a number of ATM switches, said method being adapted to handle inter-subnet communications, on a hop-by-hop basis, and characterised by the steps of said inter-subnet communications being effected using access IP/ATM nodes (AINs) and core IP/ATM nodes (CINs) of said core network; each AIN communicating with an IP/ATM host and performing IP data flow classification and labelling; an AIN, on receipt of an IP data packet from an IP/ATM host, for transmission to a destination subnet, reassembling the IP data packet; decrementing a TTL (Time To Live) timer by 1 and discarding the IP data packet if the timer reaches zero; computing the IP data packet header checksum; and checking whether the destination subnet is supported by the AIN and, if so, sending the IP data packet on that subnet. The method may be further characterised by said AIN maintaining a labelling table and a standard IP forwarding table filled by an IP routing protocol. The labelling table may include a subnet address, the outgoing VC (VPI/VCI) to reach the subnet, and a timer for the VC, the timer being started whenever IP packets are sent on the VC. The method may be further characterised by the steps of, in the event that the destination subnet is not supported by the AIN, consulting the labelling table of said AIN to determine if there is an ATM VC already in existence towards the destination subset to which the packet is to be sent; and, if an existing ATM VC is identified, said AIN hashes on the subnet address; sends the packet on the identified VC; and restarts the timer. The method may be further characterised by the steps of consulting the forwarding table of said AIN, in the event that an existing ATM VC is not identified, to find the next hop on the path to the

- 8 -

destination subnet; finding a free VCI from the VPC to the next hop (may be a CIN); sending the IP packet on a chosen VCI; creating, in the labelling table of said AIN, an entry containing the destination subnet address and the VC; and restarting the entry's timer, said timer having a lifetime value for the VC.

5

The method may be characterised by the step of using VC multiplexing (null encapsulation) of IETF RFC 1483 to encapsulate IP packets sent on said ATM VCs

The method may be characterised by the steps of purging a VC entry from the labelling table of an AIN, in the event that a corresponding VC timer times out, the purged VC being cached for a period of time before becoming available for use; forcing all downstream VCs, in the path to the destination subnet, to be timed out; and, if rerouting occurs to new routes, the affected subnet addresses found in the labelling table are assigned new VCs, based on the new routes, the old VCs becoming free after being timed out.

The method may be characterised by the steps of listening to the VPCs set up to an AIN; reassembling the cells coming from the core network into IP packets; performing computations on TTL and header checksum; and routing valid packets towards ATM hosts/routers attached to an access network.

The method may be characterised by the steps of listening to the VPCs set up to a CIN; reassembling a cell, received by the CIN on a VCI of one of said VPCs, to retrieve an IP packet; decrementing the TTL and computing the IP packet header checksum; consulting a labelling table of the CIN to determine if there is an outgoing VC already in existence to a destination subnet to which the IP packet is to be sent; and, if an existing outgoing VC is identified, the CIN sends the IP packet on that VC and restarts its timer. The method may be further characterised by the steps of consulting an IP forwarding table of the CIN, in the event that an existing outgoing VC is not identified, to find the next hop on the path towards the destination subnet; finding a free VCI from the VPC to the next hop; creating, in the labelling table of said CIN, an entry containing the destination subnet address, the incoming VC, and the outgoing VC, the free VCI chosen from the VPC to the next hop; setting the VC timer for the newly

- 9 -

created entry in the CIN's labelling table; sending the IP packet on the outgoing VC; and cross-connecting the incoming and outgoing VCs of the CIN's ATM switch, all cells thereafter received on said incoming VC being forwarded to said outgoing VC within said ATM switch. The method may be further characterised by the step of cross-connecting
5 said incoming and outgoing VCs of said ATM switch using an Ariel interface between said switch and said CIN. The method may be further characterised by said cross-connected VCs being the VCs of the incoming and outgoing Virtual Path (VP) links between the CIN's ATM switch and its two neighbouring ATM switches, one on an input side of, and the other on an output side of, said CIN's ATM switch.

The method may be characterised by said CIN periodically monitoring the cross connected VCs in order to disconnect any VC found to be idle. The method may be further characterised by using an Ariel interface between said CIN and its ATM switch to periodically monitoring the cross connected VCs. The method may be further
10 characterised by the steps of, in the event that an incoming VC is found to be inactive, the corresponding outgoing VC is purged from the CINs labelling table; the outgoing VC is put back into the labelling table as a free VC; and the incoming VC is cross-connected back to the CIN via a VP link on the CIN/ATM switch interface. Each entry in a CINs labelling table may contain a list of all incoming, merged, VCs, the outgoing VC towards
15 a destination subnet, said outgoing VC being used by the CIN during the merging of new incoming VCs and is established when a new entry is created in the labelling table, the destination subnet address, and VC timers, one for each incoming VC.

The method may be characterised by said merging of a new incoming VC with
25 the outgoing VC including the steps of each CIN listening to the VPCs set up to it, an IP packet received on a VCI of one of the VPCs being retrieved by the CIN; decrementing the TTL and computing the IP packet header checksum; consulting the CIN's labelling table to determine if there is an outgoing VC already in existence to the destination subnet to which the retrieved IP packet is to be sent; if an existing outgoing VC is
30 identified, the CIN sends the IP packet on that VC to the destination subnet, the CIN merging the incoming VC to the outgoing VC to the subnet; and, in the event that the incoming VC is not included in the list of incoming VCs in the CIN's labelling table, it is added to the table by the CIN. The method may be further characterised by the steps

- 10 -

of consulting an IP forwarding table of the CIN, in the event that an existing outgoing VC is not identified, to find the next hop on the path towards the destination subnet; finding a free VCI from the VPC to the next hop; creating, in the labelling table of said CIN, an entry containing the destination subnet address, the incoming VC (first in the VC list), the outgoing VC, the free VCI chosen from the VPC to the next hop; sending the IP packet on the outgoing VC; and merging the incoming VC to the outgoing VCs using the CIN/ATM switch interface and starting a timer for the incoming VC, all cells thereafter received on said incoming VC being forwarded to the outgoing VC within said ATM switch.

The method may be characterised by said CIN periodically monitoring the cross connected/merged VCs using the CIN/ATM switch interface in order to disconnect any VC found to be idle.

The method may be characterised by the steps of, in the event that an incoming VC is found to be inactive, the corresponding outgoing VC is purged from the CINs labelling table; the incoming VC is cross-connected back to the CIN via a VP link on the CIN/ATM switch interface; and, if the VC list becomes empty, the outgoing VC is returned as a free VC and the entry for the subnet is deleted from the CIN's labelling table.

The method may be characterised by, in the event of rerouting of an IP packet, the affected subnet addresses in the CIN's labelling table are assigned new VCs based on new routes for the packet, and the old VCs are put back as free VCs, after being timed out.

The method may be characterised by providing Quality of Service (QoS) support by setting up CBR VPCs in parallel to UBR VPCs; by each AIN assigning service classes to data flows and giving an indication of the data rate assigned to each flow using the type of service field in the IP packet header; and by a CIN, on receipt of an IP packet on a CBR VC first finding a free outgoing CBR VC towards the destination subnet; allocating the required data rate based on the type of service value; sending the IP packet on the free outgoing CBR VC; and finally, cross-connecting the incoming and

- 11 -

outgoing VCs at the CIN's ATM switch.

5 According to a further aspect of the present invention, there is provided, an IP/ATM network including an ATM transmission system as outlined in any of the preceding paragraphs, or using a method as outlined in any of the preceding paragraphs.

10 The foregoing and other features of the present invention will be better understood from the following description with reference to the accompanying drawings, in which:

15 Figure 1 diagrammatically illustrates, in the form of a block diagram, an example of a network scenario using the method (SIAM) of the present invention for the transportation of IP over ATM;

20 Figure 2 diagrammatically illustrates, in the form of a block diagram, a CIN connected to an ATM switch;

25 Figure 3 diagrammatically illustrates, in the form of a block diagram, an IP/ATM network in which SIAM is used for inter-subnet communication;

30 Figures 4 (A) and (B) diagrammatically illustrate, in block diagram form, an ATM switch before and after cross-connection of the incoming and outgoing VCs, respectively; and

Figure 5 diagrammatically illustrates, in the form of a block diagram, an ATM switch and the merging of an incoming VC with the outgoing VC.

In order to facilitate an understanding of the present invention a glossary of terms used in this patent specification is provided below:

AIN: Access IP/ATM Node

- 12 -

5

ARP: Address Resolution Protocol

ATM: Asynchronous Transfer Mode

CIN: Core IP/ATM Node

CSR: Cell Switch Router

DNS: Domain Name System

FTP: File Transfer Protocol

IETF: Internet Engineering Task Force

IGMP: Internet Group Management Protocol

IP: Internet Protocol

LAN: Local Area Network

LIS: Logical IP Subnet

MARS: Multicast Address Resolution Service

MPOA: Multi-Protocol Over ATM

NHRP: Next Hop Resolution Protocol

NHS: NHRP server

OSPF: Open Shortest Path First Protocol

PIM: Protocol Independent Multicast

20

25

30

- 13 -

P-NNI:

QoS: Quality of Service

5 RFC: Request for Comment

RSVP: Resource ReServation Protocol

SDC: Switch Diver Controller

SIAM: Simple IP/ATM Method

SVC: Switched Virtual Circuit

TTL: Time To Line *live*

UBR: Unspecified Bit Rate

UNI: User-Network Interface

20 VC: Virtual Circuit

VCI: Virtual Channel Identifier

25 VPC: Virtual Path Connection

VPI: Virtual Path Identifier

It will be seen from the subsequent description that:

30 - SIAM supports the aggregation of best effort flows which for the time being constitute the main source of traffic in the Internet;

- 14 -

- the flows are classified, based on destination subnets, which is sufficient for unicast best effort aggregation and which scales well in terms of the number of ATM VCs that are used;
- 5 - the flow classification and mapping is done at the edge of the network; and
- core nodes perform flow mapping only;

Unlike other methods, SIAM does not require any protocol for label distribution, hence a simpler architecture is achieved. The inspiration for SIAM came from the Broadband Data Service method specified and implemented, a few years ago, by Telia (see Nail Kavak, Kim Laraqui, Ala Nazari and Patrick Ernberg, The Design and Implementation of a Connectionless Broadband Data Service Protocol Over ATM, INTERNETWORKING '94, May 1994).

SIAM is based on the hop-by-hop subnet model where all communications to destinations outside the subnet go through a router. In other words, SIAM provides for inter-subnet communication within the IP/ATM network. For intra-subnet communication at the access part of the network, other IP/ATM methods are applied, for example, the Classical IP/ATM (RFC 1577), IP/PPP/ATM, etc..

Within the core network, SIAM relies on permanent VCs for best effort and multicast traffic. Whereas SVCs are used for communication at the access. SVCs also support RSVP which is mapped on ATM at the edges.

In accordance with the present invention, an ATM network will have two types of node, namely, Access IP/ATM Nodes (AINs) and Core IP/ATM Nodes (CINs), as shown in Figure 1 of the accompanying drawings which diagrammatically illustrates, in the form of a block diagram, an example network scenario using SIAM. These nodes are used for inter-subnet communication and logically segment the ATM network in the same manner as Classical IP/ATM routers. However, AINs and CINs differ from Classical IP/ATM routers in that they provide cut-through of the IP forwarding.

- 15 -

An IP host, or router, is represented in Figure 1 by a triangular block 1 and an IP/ATM host, or router, is represented by a triangular block 2. As is diagrammatically illustrated in Figure 1, the ATM network includes a core network of ATM switches and each AIN and CIN is attached to an ATM switch through an ATM User-Network Interface (ATM UNI).

The Access IP/ATM Nodes (AINs):

- perform IP flow classification and labelling, i.e. mapping IP flows to ATM VCs;
- communicate with IP/ATM hosts/routers using RFC 1577, or IP/PPP/ATM; and
- attach to the ATM network through the ATM UNI (User-Network Interface).

Non-ATM hosts can also communicate with AINs using other network technologies, for example, Ethernet.

The Core IP/ATM Nodes (CINs):

- perform routing and labelling;
- attach to the ATM network through the UNI (User-Network Interface); and
- permanent VPCs (Virtual Path Connections) are set-up between adjacent CINs.

CINs and AINs are also interconnected through VPCs. A partial mesh VPC connectivity needs to be defined for the Core.

Each flow is mapped into an ATM VC for each hop on the path towards the destination subnet. The VC for one hop between two SIAM nodes is established by selecting a free VCI from the VPC connecting the two nodes. Cut-through for one flow is achieved at a CIN via cross-connecting the flow's assigned VCs at the ATM switch to which the CIN is connected. The cross-connection is controlled by CIN and is performed

by using the Switchlet approach (see J E van der Merwe, I M Leslie, Switchlets and Dynamic Virtual ATM Networks, Integrated Network Management V, Proceedings of the Fifth IFIP/IEEE International Symposium on Integrated Network Management, May 1997).

For each CIN, a Switchlet is defined containing all the CIN's VPCs established in its ATM switch. The Switchlets of all CINs form a virtual ATM network. By using the Ariel interface the CIN cross-connects the VCI of its VPCs. The CIN's ATM interface carries both IP traffic (not yet cross-connected) as well as the Ariel interface that terminates, as shown in Figure 2, at a SDC (Switch Divider Controller). The SDC creates the Switchlet for the CIN. It can be implemented inside, or outside, the ATM switch. The same is valid for the CIN.

A CIN connected to the ATM switch and having three VPCs, set-up to three neighbours, is diagrammatically illustrated, in the form of a block diagram, in Figure 2 of the accompanying drawings.

The AINs, as described herein, do not need Switchlets because they do not perform any cross-connection at their ATM switches. However, if the host is attached via ATM and it classifies and labels flows in the same manner as the AIN, then the AIN will function as a CIN, i.e. it will perform labelling and cross-connection through an Ariel interface (see Figure 2). This is a further development of SIAM which is not addressed by this patent specification.

The Switchlet approach has been chosen because of its ability to allow each ATM switch to be shared among different control protocols such as the P-NNI and IP protocols. As is diagrammatically illustrated in Figure 2, the SDC from the Switchlet is implemented within the ATM switch but could, as stated above, be implemented outside the switch. If the current version of Ipsilon's GSMP (General Switch Management Protocol) is used instead (see P Newman et al, Ipsilon's General Switch Management Protocol Specification Version 1.1, IETF RFC 1987, August 1996), it would be necessary for the CIN to control all the resources of its switch which would then have to be used for IP traffic only. If future versions of GSMP allow the partition of switching resources

- 17 -

among different control architectures, then SIAM could use GSMP instead of Ariel.

As stated above, SIAM is used for inter-subnet communication. Intra-subnet communication at the access relies on other IP/ATM methods, for example, the Classical IP/ATM method (RFC 1577), IP/PPP/ATM etc.. As is diagrammatically illustrated in Figure 3 of the accompanying drawings, in the form of a blocked diagram, one of the IP routing protocols will run in the Core Network, for example, the OSPF protocol can run on the pre-configured VPCs using the point-to-multipoint interface, or the Designated Router approach.

The manner in which SIAM nodes, AIN and CIN, operate will now be described with reference to the accompanying drawings.

In the case of AINs, each AIN maintains a labelling table and a standard IP forwarding table filled by the IP routing protocol. Each entry in the labelling table includes a subnet address, the outgoing VC (VPI/VCI) to reach that subnet and a timer for the VC. The timer is restarted whenever IP packets are sent on the VC. Therefore, entries in the labelling table are maintained on a flow basis (using timeouts) rather than at a packet level.

Thus, in operation, i.e. access side to core network/access, SIAM includes the following steps:

1. An IP packet received from the IP/ATM host is reassembled. As illustrated in Figure 1, packets may also arrive through non ATM interfaces;
2. TTL (Time To Live) is decremented by 1 and the packet is discarded if the result reaches zero. The IP header checksum is computed;
3. A check is made to determine whether the destination subnet is also supported by the AIN. If it is supported by AIN, the packet is sent to that subnet. If not so supported, the method progresses to step 4;

- 18 -

4. In the event that the destination subnet is not supported by the AIN, the labelling table is consulted to find if there is an ATM VC already in existence towards the same destination subnet. If an existing ATM VC is identified, the AIN hashes on the subnet address, sends the packet on the identified VC and restarts its timer. If an existing ATM VC cannot be identified, the method progresses to step 5;

5. In the event that an existing ATM VC is not identified, the IP forwarding table is consulted to find the next hop on the path to the destination subnet. The AIN then finds a free VCI from the VPC to that hop (for example, CIN). The packet will be sent on the VC (i.e. the chosen VCI from the VPC to the next hop). An entry in the labelling table is created containing the destination subnet address and the VC. The entry's timer is also started with a VC lifetime value (parameterized).

The VC multiplexing (null encapsulation) of RFC 1483 is used to encapsulate IP packets sent on the ATM VCs.

When a VC timer times out, the corresponding VC entry is purged from the labelling table. However, the VC does not become available immediately, but is cached for a while. By doing so, all the involved VCs at the path downstream are forced to time out too.

If rerouting occurs, the affected subnet addresses found in the labelling table are assigned new VCs, based on the new routes, and the old VCs become free after being timed out.

As for the core network side to access:

1. The AIN listens to the VPCs set up to it;
2. Cells coming from the core network are reassembled into IP packets;
3. Computations are performed on TTL and header checksum; and

- 19 -

4. Valid packets are then routed towards the hosts/routers attached to the access network.

5 In the case of CINs, each CIN also maintains a labelling table and a standard IP forwarding table. Each entry of the labelling table includes a subnet address, an outgoing ATM VC to reach that subnet and the incoming VC on which the packets arrive. A VC timer is set for each entry and it is restarted whenever packets are received on the incoming VC. When an incoming VC has not been active during its lifetime, its corresponding outgoing VC will become free and will be at the disposal of other flows. As with the AIN, the VC does not become free directly, thus forcing the VCs on the downstream hops to time out too.

Each CIN operates as follows:

1. The CIN listens to the VPCs set up to it. An IP packet received on a VCI of one of these VPCs is reassembled retrieving the IP packet;
2. The TTL is decremented and IP header is calculated;
3. The labelling table is consulted to determine (for example, by hashing on the incoming VC) if there is an outgoing VC already in existence to the same destination subnet. This indicates that the incoming and outgoing VCs have not yet been cross-connected (see Figure 4 (A) of the accompanying drawings). If an existing outgoing VC is identified, the CIN sends the packet on the outgoing VC and restarts its timer. If an existing outgoing VC cannot be identified, the method progresses to step 4;
4. In the event that an outgoing VC cannot be identified, the IP forwarding table is consulted to find the next hop on the path toward the destination subnet. The CIN then finds a free VCI from the VPC to that hop. An entry in the labelling table is created containing the destination subnet address, the incoming VC and the outgoing VC, i.e. the chosen VCI on the VPC to the next hop. It then sets the

- 20 -

5 outgoing VC timer for the newly created entry. The CIN sends the packet on the outgoing VC. It then cross-connects the incoming and the outgoing VCs through use of the Ariel interface. The VCs of the incoming and outgoing VP links (between the CIN's switch and its two neighbours) are actually cross- connected (see Figure 4 (B) of the accompanying drawings). Thereafter, all cells received on the incoming VC are forwarded to the outgoing VC directly in the ATM switch, as shown in Figure 4 (B).

10 By using the Ariel interface (see Figure 2), the CIN periodically monitors its cross-connected VCs in order to disconnect the idle ones. If it finds an incoming VC that has not been active during its lifetime, the corresponding VC entry is purged from the labelling table. The outgoing VC is put back in the labelling table as a free VC and finally the incoming VC is cross-connected back to the CIN (the VP link on the CIN interface - see Figure 4 (A) of the accompanying drawings). If cells are received on one incoming VC, its timer is restarted.

15 If rerouting occurs, the affected subnet addresses found in the labelling table are assigned new VCs based on the new routes and the old VCs become free after being timed out.

20 As stated above, CINs perform flow classification based on source and destination subnet addresses. Since the AINs normally serve more than one access subnet and aggregate the traffic from all their access subnets into one VC toward one destination subnet, the classification is based on a number of access subnets and one destination subnet.

25 CINs can also be defined to perform flow classification based on destination subnet addresses only, achieving VC conservation and a smaller number of entries in a CIN's labelling table. For this purpose, the ATM switch must perform VC merging, i.e. multipoint-to-point, which is not, at the present time, a standard function of an ATM switch. It thus becomes a "packet switch". The operation of CINs also becomes rather complex.

- 21 -

The merging of an incoming CV of an ATM switch with the outgoing CV is diagrammatically illustrated, in the form of a block diagram, in Figure 5 of the accompanying drawings.

5 Each entry of the labelling table of the CIN contains:

- a list of all the incoming VCs, i.e. the merged VCs;
- the outgoing VC toward the destination subnet;
- the destination subnet address; and
- VC timers, one for each incoming VC.

The outgoing VC is used by the CIN during the merging process for new incoming VCs and is established when a new entry is created in the labelling table.

Each CIN operates as follows:

- 20 1. The CIN listens to the Virtual Path Connections (VPCs) set-up to it; an IP packet received on a VCI of one of these VPCs is retrieved;
2. The TTL is decremented and the IP header checksum is calculated;
- 25 3. The labelling table of the CIN is consulted to determine if there is an outgoing VC already in existence to the same destination subnet to which the IP packet has to be sent.
- 30 4. If an existing outgoing VC is identified, the CIN sends the IP packet on the identified outgoing VC for the subnet and, by using the Ariel interface, the CIN merges the incoming VC to the outgoing VC to the subnet. The CIN also adds the incoming VC (if it is not already there) to the list of incoming VCs and starts a timer for it.

- 22 -

5. In the event that there is no entry for the destination subnet, i.e. an existing outgoing VC cannot be identified, the method progresses to step 6;

5 6. In the absence of an existing outgoing VC, the CIN consults the IP forwarding table to find the next hop on the path to the destination subnet. The CIN then finds a free VCI from the VPC to that hop and an entry in the labelling table is created containing:

- 10
- the destination subnet address;
 - the incoming VC (first in the VC list); and
 - the outgoing VC (the chosen VCI on the VPC to the next hop).

15 7. The CIN then sends the packet on the outgoing VC. By using the Ariel interface, the CIN merges the incoming VC to the outgoing VC. It also starts a timer for the incoming VC. After this, all cells received on the incoming VC are forwarded to the outgoing VC directly in the ATM switch, as shown in Figure 5.

20 The CIN periodically monitors the incoming VCs in order to disconnect the idle ones. If it finds a VC that has not been active during its lifetime, the VC is purged from the VC list and is cross-connected back to the CIN. If the VC list becomes empty, the outgoing VC is returned as a free one and the entry for the subnet is deleted from the labelling table. If cells are received on one incoming VC, its timer is restarted.

25

If rerouting occurs, the affected subnet addresses found in the labelling table are assigned new VCs based on the new routes and the old VCs are put back as free ones after being timed out. The incoming VCs will also be merged to the new outgoing VC.

30 Quality of Service (QoS) support for the system could be based on the use of RSVP (see R Braden, L Zhan, S Berson, S Herzog, S Jamin, Resource ReReservation Protocol (RSVP) - Version 1 Functional Specification, draft-ietf-rsvp-spec-15.txt, May 27,

1997).

The IP-ATM service mapping, i.e. choosing the ATM service categories for the integrated service classes, and VC management, i.e. mapping RSVP sessions to ATM VCs, are performed at the AINs. RSVP merging is also implemented at the AINs. On the other hand, CINs do not support RSVP. The RSVP messages are transported and treated within the Core Network as best effort traffic. In other words, CINs are non-RSVP routers. Used this way, the scalability problem of RSVP in Core Networks is avoided.

The manner in which RSVP is implemented on SIAM is not addressed, in great detail, by this patent specification. However, for the purpose of the present invention, it should be noted that, when an AIN receives a Resv message from a host, it adds its ATM address to the message and sends it upstream toward the source. When the source's AIN receives the Resv message, it performs VC management by setting-up an SVC to the receiver's AIN using the ATM address carried by the Resv message.

As an interim solution, a simplified collection of QoS can be supported. Besides best effort, a guaranteed service class is also provided. The best effort is mapped onto the ATM UBR (Unspecified Bit Rate) extended with some packet discard mechanism, for example, early packet discard. The manner in which best effort is implemented with SIAM is outlined in the preceding description.

For the guaranteed service, CBR VPCs are set-up in parallel to UBR VPCs. AINs assign service classes to flows according to some policy control that may be based on source subnets. In case of the guaranteed service, AINs also indicate the data rate assigned to each flow by using the type of service field of the IP packet header.

A CIN, on receipt of a packet on a VC of the type CBR:

- first finds a free outgoing CBR VC toward the destination;
- allocates the required data rate based on the type of service value;

- 24 -

- sends the packet on the outgoing VC; and
- finally, cross-connects the incoming and outgoing VCs at the ATM switch.

5

As for IP multicasting, the known approaches for the access side can be implemented. In the case of Classical IP/ATM in the access, intra-LIS multicasting can be based on the MARS approach (see G Armitage, Support for Multicast over UNI 3.0/3.1 based ATM Networks, RFC 2022, November 1996). In case of using IP/PPP/ATM or other known methods, the AIN functions as a multicast router supporting IGMP (Internet Group Management Protocol).

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
105
110
115
120
125
130
135
140
145
150
155
160
165
170
175
180
185
190
195
200
205
210
215
220
225
230
235
240
245
250
255
260
265
270
275
280
285
290
295
300
305
310
315
320
325
330
335
340
345
350
355
360
365
370
375
380
385
390
395
400
405
410
415
420
425
430
435
440
445
450
455
460
465
470
475
480
485
490
495
500
505
510
515
520
525
530
535
540
545
550
555
560
565
570
575
580
585
590
595
600
605
610
615
620
625
630
635
640
645
650
655
660
665
670
675
680
685
690
695
700
705
710
715
720
725
730
735
740
745
750
755
760
765
770
775
780
785
790
795
800
805
810
815
820
825
830
835
840
845
850
855
860
865
870
875
880
885
890
895
900
905
910
915
920
925
930
935
940
945
950
955
960
965
970
975
980
985
990
995

Some CINs function as multicast routers and run PIM (Protocol Independent Multicast) of RFC 2117 (see D Estrin et al, Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification, RFC 2117, June 1997). PIM is also used for inter-LIS multicast communication where the AINs become multicast routers member of the multicast groups found at the access. As a first step, multicasting within the Core Network is not supported by short-cut VCs which are introduced later. Furthermore, the multicast routers are interconnected via permanent VCs.

20

As for interworking with ATM, conventional routers and hosts, host communicate with their AINs using existing IP/ATM methods implying that SIAM is transparent for the hosts.

25

The AINs can also communicate with hosts and routers based on legacy networks, for example, Ethernet.

30

AINs and CINs interconnect with ATM switches using the UNI interface. AINs are IP/ATM-routers extended with flow aggregation capabilities. The same applies to CINs that implement both flow aggregation and VC cross-connection based on, for example, the Switchlet approach that enables the ATM switches to be used also for other applications other than IP.

- 25 -

It will be seen from the foregoing description of the present invention that SIAM:

- provides an IP/ATM mechanism that requires minor changes to the current IP/ATM infrastructure;
- is independent from the IP routing protocol used;
- does not require the use of any label distribution protocol which characterizes the existing methods, for example, IP switching, Tag switching, and CSR;
- provides low latency transport of IP best effort traffic; and
- can support QoS and multicasting.

It will also be seen from the foregoing description that, within the Core Network, where there is a high level of traffic aggregation and less dynamic behaviour than the edges, SIAM makes use of permanent VPCs that are quite sufficient to interconnect a limited number of CINs but a high number of access subnets. As a result, the connection set-up time is minimized and there is no need for any time-consuming address resolution. At the access, where the largest number of nodes reside, for example, hosts and AINs, communication is less static and SVCs are, therefore, used to make efficient use of access network resources. SVCs can also support RSVP where the IP-ATM service mapping and VC management are performed at AINs.

SIAM provides traffic aggregation based on destination subnet addresses and hence achieving a scalable architecture. It is both traffic driven as well as topology driven meaning that ATM VCs are set-up when they are first needed, but that VPs between CINs are set up based on topological considerations.

It will be readily understood by persons skilled in the art that SIAM can be used to build efficient public broadband IP/ATM networks.

SIAM follows the current trend in the industry focussing on the cut-through of the

- 26 -

IP layer processing. However, SIAM has a simpler architecture because it does not require the use of any label distribution protocol which, as stated above, characterizes existing methods.

5

100
99
98
97
96
95
94
93
92
91
90
89
88
87
86
85
84
83
82
81
80
79
78
77
76
75
74
73
72
71
70
69
68
67
66
65
64
63
62
61
60
59
58
57
56
55
54
53
52
51
50
49
48
47
46
45
44
43
42
41
40
39
38
37
36
35
34
33
32
31
30
29
28
27
26
25
24
23
22
21
20
19
18
17
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1

20

25

30